



Intel® Xeon™ Processor

Specification Update

September 2003

Notice: The Intel® Xeon™ processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Document Number: 249678-030



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS. INTEL MAY MAKE CHANGES TO SPECIFICATIONS AND PRODUCT DESCRIPTIONS AT ANY TIME, WITHOUT NOTICE.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Xeon™ Processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel, Intel Xeon, Pentium, Celeron and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2001-2003, Intel Corporation



Contents

Revision History 4

Preface 7

Errata 22

Specification Changes..... 46

Specification Clarifications 48

Revision History

Version	Description	Date
-001	Initial Release.	May 2001
-002	Added errata P27-P28.	June 2001
-003	Updated erratum P25.	June 2001
-004	Added RCPPS, RCPSS, RSQRTPS and RSQRTSS instruction specification clarification.	July 2001
-005	Added errata P29-P32. Added Unused outputs specification clarifications.	August 2001
-006	Added errata P33-34. Production mark update to include 2D matrix	September 2001
-007	Added 2 GHz frequency, D0 stepping and FC-BGA package information Added DP Platform Population Matrix Updated Summary of Errata table with applicable "Fixed" & "No Fix" errata plans Updated Errata P32	October 2001
-008	Added Erratum P35 Updated Summary of Errata Table with latest errata status Added Documentation Changes P1-P5	November 2001
-009	Added errata P36 and P37. Added Documentation Change P6. Updated Summary Tables with latest errata and documentation changes.	December 2001
-010	Added Documentation Changes P7-P11. Added Errata P38-P44. Added new processor with Processor Signature=0F24h B0 Step. Added new Sspec processors to Processor ID and DP Matrix tables. Added new package markings(Figs 3,4) for 512KB Cache processor.	January 2002
-011	Corrected 3 mislabeled S-Spec Table parts.	January 2002
-012	Added Erratum P45 to Processor ID Table.	February 2002
-013	Added Errata P46 and P47. Added Documentation Changes P1-P6. Minor changes to "Fix" descriptions.	April 2002
-014	Added PWRGOOD Specification Change. Added Errata P48. Updated Errata P11. Added Specification Changes P1-P3. Added Documentation Changes P1-P3. Added new S-specs SL687 and SL65T to Processor ID table. Updated datasheet name to 2.40GHz	May 2002
-015	Added Erratum P49. Updated Errata P24, P40 status Added Document Changes P1-P2	June 2002



Version	Description	Date
-016	Added new Erratum P50. Added new Doc Changes P3-P12. Edited Summary of Errata Table erratum P40 to PlanFix. Minor edits to processor markings.	July 2002
-017	Edited DP Matrix table. Updated the Summary of Errata table w C1 Step info.	August 2002
-018	Added errata P51, P52. Edited erratum P14. Added Doc Changes P3-P24. Added Spec Clarification P1. Added new C1 S-specs to Processor ID Info table.	September 2002
-019	Edited erratum P35. Removed erratum previously numbered P52 as not applicable. Added new erratum P52. Added Doc Changes P25-P32.	October 2002
-020	Added erratum P53. Added new Spec Change P1. Added new Spec Clarification P2. Deleted old Doc Changes. Deleted old Spec Changes Added new Processor (Intel® Xeon™ Processor with 533 MHz Front Side Bus Added new processor with Processor Signature=0F27h C1 Step. Added New S-specs to the processor ID table All References to CPUID are now renamed Processor Signature	November 2002
-021	Added reference to IA32 Software Developers Manual	December 2002
-022	Added Erratum 054-055	January 2003
-023	Added new Errata P56, P57. Added boxed processors S-Specs for 400 MHz FSB with the frequencies: 1.80, 2, 2.20 and 2.40 GHz.	February 2003
-024	Added new Errata P58 and P59. Updated previous Errata fixing status, P55. Updated the name of the datasheet reference. Added S Spec #SL6RR, SL6GH.	March 2003
-025	Updated previous Errata P55 and P58	April 2003
-026	Added Specification Clarification P1-P4 Added Low Voltage Intel Xeon Specifications	May 2003
-027	Added Erratum P60 Updated Sections with Low Voltage Intel Xeon Processor Datasheets Updated Erratum P53	June 2003
-028	Added new Processor Intel® Xeon™ Processor with 1-MB L3 Cache with Processor Signature=0F25H (M0 Stepping). Added New S-specs to the processor ID table for 0F25H (M0 Stepping) and 0F29h (D1 Stepping) Added new processor with Processor Signature=0F29H (D1 Stepping) Updated DP Platform Population Matrix for the Intel® Xeon™	July 2003

Version	Description	Date
	Processor to include 0F25H and 0F29H Removed Spec Clarification P3 Updated Erratum P30 – Title Correction	
-029	Updated Errata P54 and P56	August 2003
-030	Added erratum P64 Updated Erratum P9, P22, P39	September 2003



Preface

This document is an update to the specifications contained in the documents listed in the following Affected Documents/Related Documents table. It is a compilation of device and document errata and specification clarifications and changes, and is intended for hardware system manufacturers and for software developers of applications, operating system, and tools.

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents/Related Documents

Document Title	Document Number
<i>Intel® Xeon™ Processor at 1.40 GHz, 1.50 GHz, 1.70 and 2 GHz datasheet</i>	249665
<i>Intel® Xeon™ Processor with 512 KB L2 Cache at 1.80 GHz to 2.80 GHz Datasheet</i>	298642
<i>Intel® Xeon™ Processor with 533 MHz Front Side Bus at 2 GHz to 3.06 GHz Datasheet</i>	252135
<i>Low Voltage Intel® Xeon™ Processor at 1.6 GHz and 2 GHz Datasheet</i>	273766
<i>IA-32 Intel® Architecture Software Developer's Manual, Volumes 1, 2 and 3</i>	245470, 245471, and 245472, respectively

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L2 cache size, package type, etc., as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. Errata may cause the Intel Xeon processor's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Changes are modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

Intel® Xeon™ Processor Markings, 256-KB Cache (603-pin interposer INT-mPGA package)

Figure 1. Top Side Processor Marking

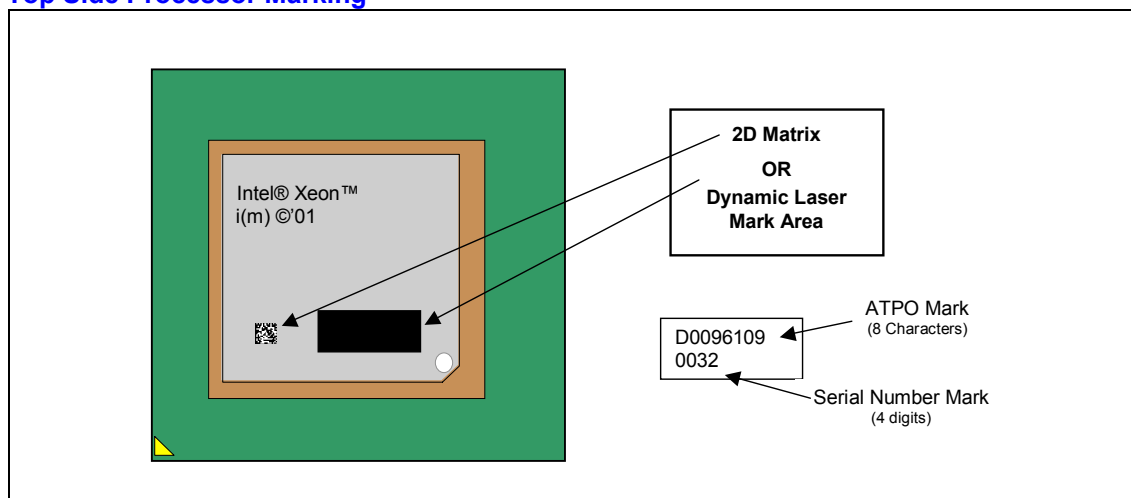
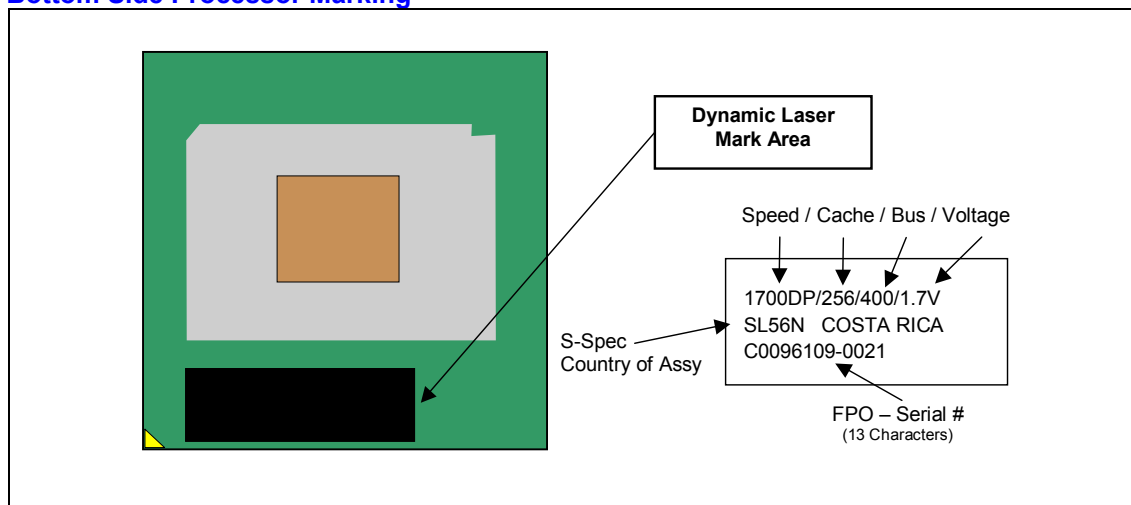


Figure 2. Bottom Side Processor Marking



Intel® Xeon™ Processor , 512-KB Cache, 400 and 533 MHz FSB Markings, (603-pin interposer INT-mPGA package and 604-pin Fc-mPGA2 Package)

Figure 3. Top Side Processor Marking




Figure 4. Bottom Side Processor Marking

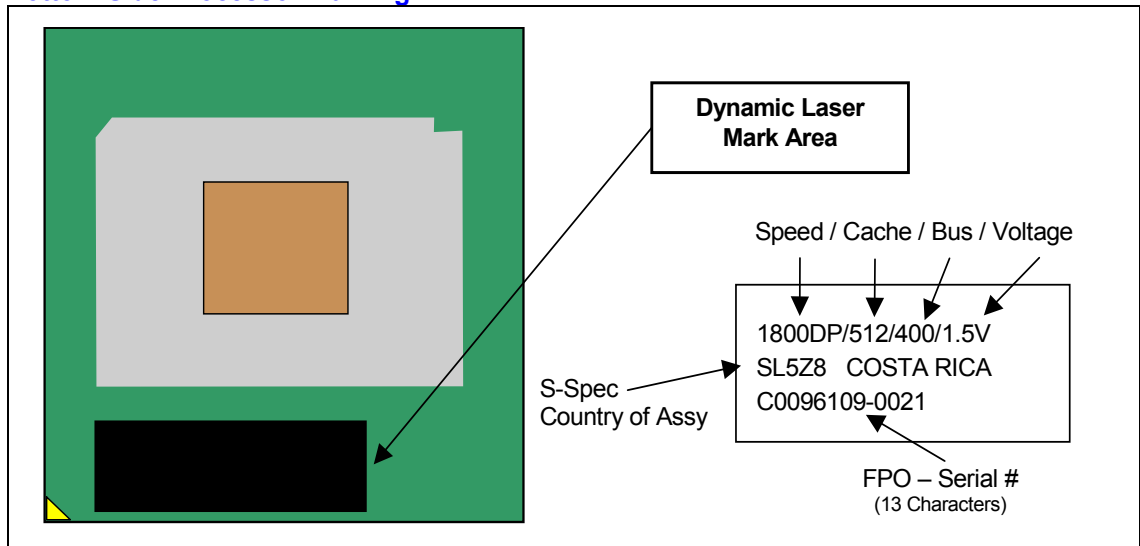


Figure 5. Example of Production Mark – Top View

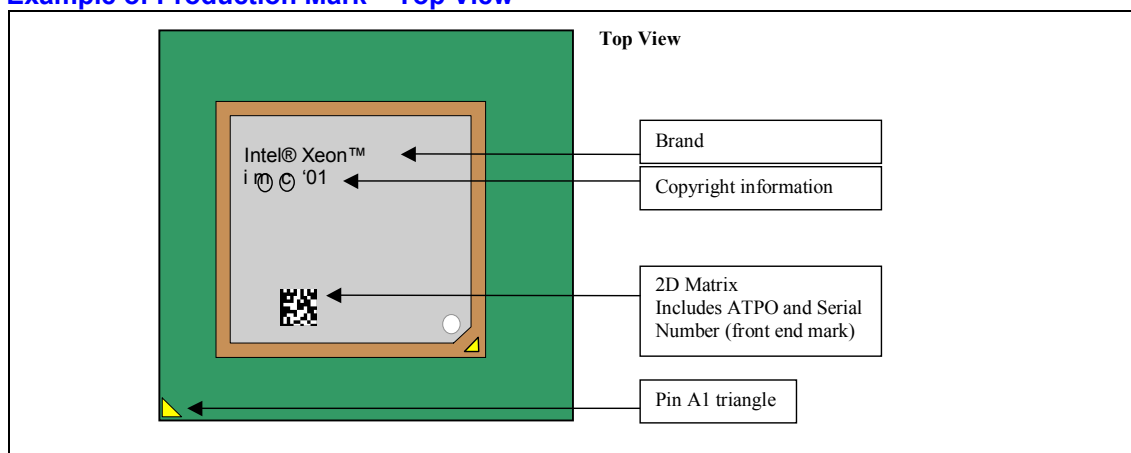
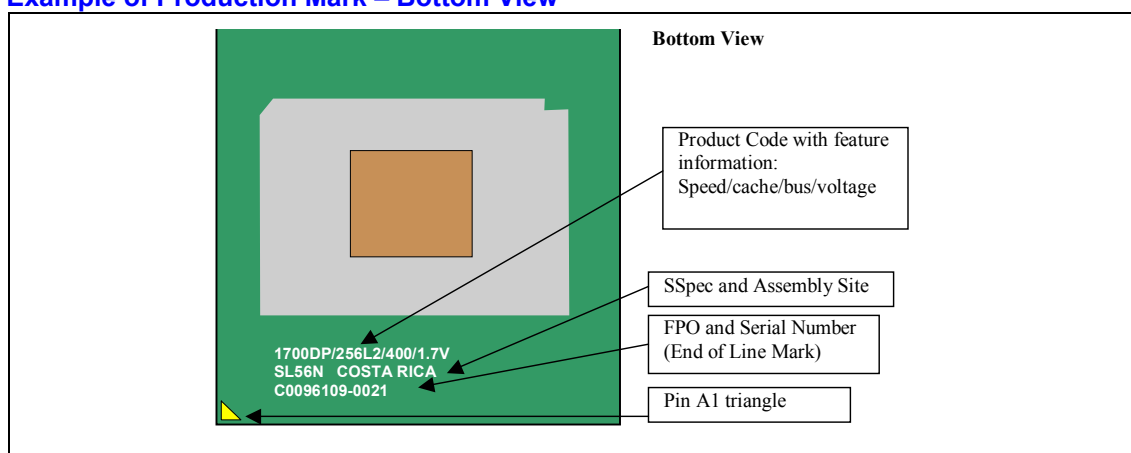


Figure 6. Example of Production Mark – Bottom View



Identification Information

The Intel Xeon processor can be identified by the following values:

Family ¹	Model ²	Brand ID ³
1111	0000	00001110
1111	0001	00001110
1111	0010	00001011

The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the Processor Signature instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.

The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the Processor Signature 2 instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

The Brand ID corresponds to bits [7:0] of the EBX register after the Processor Signature instruction is executed with a 1 in the EAX register.

Cache and TLB descriptor parameters are provided in the EAX, EBX, ECX and EDX registers after the Processor Signature instruction is executed with a 2 in the EAX register. Please refer to the Intel



Processor Identification and the Processor Signature Instruction Application Note (AP-485) for further information on the Processor Signature instruction.

Table 1. Intel® Xeon™ Processor Identification and Package Information

S-Spec Number	Core Stepping	Processor Signature	Speed Core/Front Side Bus (GHz/MHz)	L2 Size (Kbytes)	L3 Size (Kbytes)	Processor Interposer Revision	Package and Revision	Notes
SL4WX	C1	0F0Ah	1.40/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 4
SL56G	C1	0F0Ah	1.40/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 2, 4
SL4WY	C1	0F0Ah	1.50/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 4
SL4ZT	C1	0F0Ah	1.50/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 2, 4
SL56N	C1	0F0Ah	1.70/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 4
SL56H	C1	0F0Ah	1.70/400	256-KB		B0	603-pin micro-PGA interposer with 31 mm OLGA rev 2.0	1, 2, 4
SL5TD	D0	0F12h	1.50/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package	1, 4
SL5U6	D0	0F12h	1.50/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package	1, 2, 3, 4
SL5TE	D0	0F12h	1.70/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package.0	1,4
SL5U7	D0	0F12h	1.70/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package	1, 2, 3, 4
SL5TH	D0	0F12h	2/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package 0	1, 3, 4
SL5U8	D0	0F12h	2/400	256-KB		C0	603-pin micro-PGA interposer with 31 mm FC-BGA package	1, 2,3,4
SL5Z8	B0	0F24h	1.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1
SL622	B0	0F24h	1.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1, 2
SL5Z9	B0	0F24h	2/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1

S-Spec Number	Core Stepping	Processor Signature	Speed Core/Front Side Bus (GHz/MHz)	L2 Size (Kbytes)	L3 Size (Kbytes)	Processor Interposer Revision	Package and Revision	Notes
SL623	B0	0F24h	2/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1, 2
SL5ZA	B0	0F24h	2.20/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1
SL624	B0	0F24h	2.20/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1, 2
SL687	B0	0F24h	2.40/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	1,2
SL65T	B0	0F24h	2.40/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6EL	C1	0F27H	1.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6JX	C1	0F27H	1.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6EM	C1	0F27H	2/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6JY	C1	0F27H	2/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	2
SL6EN	C1	0F27H	2.20/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6JZ	C1	0F27H	2.20/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	2
SL6EP	C1	0F27H	2.40/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6K2	C1	0F27H	2.40/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	2
SL6EQ	C1	0F27H	2.60/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6K3	C1	0F27H	2.60/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	2
SL6M7	C1	0F27H	2.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	
SL6MS	C1	0F27H	2.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	2
SL6NP	C1	0F27H	2/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,5



S-Spec Number	Core Stepping	Processor Signature	Speed Core/Front Side Bus (GHz/MHz)	L2 Size (Kbytes)	L3 Size (Kbytes)	Processor Interposer Revision	Package and Revision	Notes
SL6NQ	C1	0F27H	2.40/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,5
SL6NR	C1	0F27H	2.66/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,5
SL6NS	C1	0F27H	2.8/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,5
SL6RQ	C1	0F27H	2/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6GD	C1	0F27H	2.40/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6GF	C1	0F27H	2.66/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6GG	C1	0F27H	2.8/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6GH	C1	0F27H	3.06/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5, 6
SL6RR	C1	0F27H	3.06/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,6
SL6GV	C1	0F27H	1.60/400	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	7
SL6W3	D1	0F29H	1.80/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6W6	D1	0F29H	2/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6W7	D1	0F29H	2.20/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6W8	D1	0F29H	2.40/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6W9	D1	0F29H	2.60/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6WA	D1	0F29H	2.8/0400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3
SL6WB	D1	0F29H	3/400	512-KB		01	603-pin micro-PGA interposer with 35 mm FC-BGA package	3,6
SL6XK	D1	0F29H	1.6/400	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	7

S-Spec Number	Core Stepping	Processor Signature	Speed Core/Front Side Bus (GHz/MHz)	L2 Size (Kbytes)	L3 Size (Kbytes)	Processor Interposer Revision	Package and Revision	Notes
SL6XL	D1	0F29H	2/400	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	7
SL6VK	D1	0F29H	2/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6VL	D1	0F29H	2.40/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6VM	D1	0F29H	2.66/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6VN	D1	0F29H	2.8/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5
SL6VP	D1	0F29H	3.06/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	5,6
SL73K	M0	0F25H	2/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2
SL72C								
SL73L	M0	0F25H	2.40/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2
SL72D								
SL73M	M0	0F25H	2.66/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2
SL72E								
SL73N	M0	0F25H	2.80/533	512-KB		01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2
SL72F								
SL73P	M0	0F25H	3.06/533	512-KB	1-MB	01	604-pin micro-PGA interposer with 42.5 mm FC-PGA2 package	2,6,8
SL72G								6,8

Notes

1. The Intel Xeon processor listed here is installed onto a micro pin grid array (mPGA) interposer. The overall processor package is called INT-mPGA.
2. These parts are Intel boxed processors.
3. FC-BGA packaging maintains form, fit, and functionality when compared to OLGA packaging. Users may notice a color change.
4. These parts require the inputs from A20M#, IGNNE#, LINT[1]/NMI and LINT[0]/INTR pins during RESET to set the correct core to bus frequency ratio.
5. These parts are the Intel® Xeon™ Processor with 533 MHz Front Side Bus
6. These parts have a VID of 1.525V
7. These parts are the Low Voltage Intel Xeon Processor
8. These parts are the Intel® Xeon™ Processor with 1-MB L3 Cache



Mixed Steppings in DP Systems

Intel Corporation fully supports mixed steppings of Intel® Xeon™ processors. The following list and processor matrix describes the requirements to support mixed steppings:

- Mixed steppings are only supported with processors that have identical family numbers as indicated by the Processor Signature instruction. The Intel® Xeon™ processor is available with two different Model numbers as indicated by the Processor Signature. Please refer to the “DP Platform Population Matrix for the Intel® Xeon™ Processor” for details regarding inclusion of processors with mixed Processor Signature/Core steppings.
- While Intel has done nothing to specifically prevent processors operating at differing frequencies from functioning within a multiprocessor system, there may be uncharacterized errata that exist in such configurations. Intel does not support such configurations. In mixed stepping systems, all processors must operate at identical frequencies (i.e., the highest frequency rating commonly supported by all processors).
- While there are no known issues associated with the mixing of processors with differing cache sizes in a multiprocessor system, and Intel has done nothing to specifically prevent such system configurations from operating, Intel does not support such configurations since there may be uncharacterized errata that exist. In mixed stepping systems, all processors must be of the same cache size.
- While Intel believes that certain customers may wish to perform validation of system configurations with mixed frequency or cache sizes, and that those efforts are an acceptable option to our customers, customers would be fully responsible for the validation of such configurations.
- Intel requires that the proper microcode update be loaded on each processor operating in a multiprocessor system. Any processor that does not have the proper microcode update loaded is considered by Intel to be operating out of specification.
- The workarounds identified in this and following specification updates must be properly applied to each processor in the system. Certain errata are specific to the multiprocessor environment and are identified in the *Mixed Stepping Processor Matrix* found at the end of this section. Errata for all processor steppings will affect system performance if not properly worked around. Also see the Intel® Xeon™ Processor Identification and Package Information section for additional details on which processors are affected by specific errata.
- In mixed stepping systems, the processor with the lowest feature-set, as determined by the Processor Signature Feature Bytes, must be the Bootstrap Processor (BSP). In the event of a tie in feature-set, the tie should be resolved by selecting the BSP as the processor with the lowest stepping as determined by the Processor Signature instruction.

In the following processor matrix, “**NI**” indicates that there are currently no known issues associated with mixing these steppings. A number indicates that a known issue has been identified as listed in the table following the matrix. “**X**” indicates the processors cannot be mixed. A dual processor system using mixed processor steppings must assure that errata are addressed appropriately for each processor.

Table 2. DP Platform Population Matrix for the Intel® Xeon™ Processor²

Processor Signature/Core Stepping	0F0Ah/C1	0F12h/D0	0F24h/B0	0F27h/C1	0F29h/D1	0F25h/M0 ³	0F25h/M0 ⁴
0F0Ah/C1	NI	Note 1	X	X	X	X	X
0F12h/D0	Note 1	NI	X	X	X	X	X
0F24h/B0	X	X	NI	NI	NI	NI	X
0F27h/C1	X	X	NI	NI	NI	NI	X
0F29h/D1	X	X	NI	NI	NI	NI	X
0F25h/M0	X	X	NI	NI	NI	NI	NI

Notes:

1. Some of these processors are affected by errata, which may affect the features an MP system is able to support. See the *Intel® Xeon™ Processor Identification and Package Information* table for details on which processors are affected by these errata.
2. This Matrix also applies to the Intel® Xeon™ Processor with 533 MHz Front Side Bus, Low Voltage Intel Xeon Processor, and Intel Xeon Processor with 1-MB L3 cache
3. This only applies to 0F25h stepping with no L3 cache
4. This only applies to 0F25h stepping with no L3 cache



Summary Tables of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed MCH steppings. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

Codes Used in Summary Table

X:	Erratum, Specification Change or Clarification that applies to the given processor stepping.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Doc:	Document change or update that will be implemented.
PlanFix:	This erratum may be fixed in a future of the product.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
PKG:	This column refers to errata on the Intel Xeon processor substrate
AP	APIC-related erratum
Shaded:	This item is either new or modified from the previous version of the document.

Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

- A = Intel® Pentium® II processor
- B = Mobile Intel® Pentium® II processor
- C = Intel® Celeron® processor
- D = Intel® Pentium® II Xeon™ processor
- E = Intel® Pentium® III processor
- G = Intel® Pentium® III Xeon™ processor
- H = Mobile Intel® Celeron® processor at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz
- K = Mobile Intel® Pentium® III Processor - M
- M = Mobile Intel® Celeron® processor
- N = Intel® Pentium® 4 processor
- O = Intel® Xeon™ processor MP and for Intel® Xeon™ processor MP with up to 2-MB L3 cache on 0.13-micron process
- P = Intel® Xeon™ processor , Intel® Xeon™ processor with 512-KB L2 Cache, Intel® Xeon™ processor with 533 MHz Front Side Bus, Low Voltage Intel Xeon Processor , and Intel® Xeon™ processor with 1-MB L3 cache
- T = Mobile Intel® Pentium® 4 processor
- V = Mobile Intel® Celeron® processor on .13 Micron Process in Micro-FCPGA Package
- Y = Intel® Pentium® M processor

Note: The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.

SUMMARY OF ERRATA								
NO.	C1/ 0F0Ah	D0/ 0F12h	B0/ 0F24h	C1/ 0F27h	D1/ 0F29h	M0/ 0F25h	PLANS	ERRATA
P1	X	X					Fixed	UC Code in same line as WriteBack (WB) data may lead to data corruption
P2	X	X	X	X	X	X	No Fix	Transaction is not retried after BINIT#
P3	X	X	X	X	X	X	No Fix	Invalid opcode 0FFFH requires a ModRM byte
P4	X	X	X	X	X	X	No Fix	When in No-Fill Mode (CR0.CD=1) the memory type of large (PSE-4M and PAE-2M) pages are wrongly forced to uncacheable
P5	X	X	X	X	X	X	No Fix	Processor may hang due to Speculative Page Walks to Non-Existent System Memory
P6	X	X					Fixed	Writing a performance counter may result in an incorrect counter value
P7	X	X					Fixed	Performance Counter May Contain Incorrect Value After Being Stopped
P8	X						Fixed	REP MOV instruction with overlapping source and destination may result in data corruption
P9	X	X	X	X	X	X	No Fix	Memory type of the load lock different from its corresponding store unlock
P10	X	X	X	X	X	X	No Fix	Machine check architecture error reporting and recovery may not work as expected
P11	X	X	X	X	X	X	No Fix	Debug mechanisms may not function as expected
P12	X						Fixed	Processor may live-lock if PDEs or PTEs are in UC space
P13	X						Fixed	Thermal status log bit may not be set when the Thermal Control Circuit is Active
P14	X	X					Fixed	Processor may timeout waiting for a device to respond after 0.67 seconds
P15	X	X	X	X	X	X	No Fix	Cascading of performance counters does not work correctly when forced overflow is enabled
P16	X	X	X	X	X	X	No Fix	EMON event counting of X87 loads may not work as expected
P17	X	X					Fixed	Simultaneous code breakpoint and uncorrectable error results in a processor hang
P18	X	X					Fixed	Software controlled clock modulation using a 12.5% or 25% duty cycle may cause the processor to hang
P19	X						Fixed	RFO with ECC error may result in data corruption
P20	X						Fixed	Speculative page fault may cause livelock
P21	X						Fixed	PAT index MSB may be calculated incorrectly
P22	X	X	X	X	X	X	No Fix	System bus interrupt messages without data and which receive a HardFailure response may hang the processor



SUMMARY OF ERRATA								
NO.	C1/ 0F0Ah	D0/ 0F12h	B0/ 0F24h	C1/ 0F27h	D1/ 0F29h	M0/ 0F25h	PLANS	ERRATA
P23	X	X					Fixed	SQRTPD and SQRTSD May Return QNaN Indefinite Instead of Negative Zero
P24	X	X	X				Fixed	Bus Invalidate Line Request that return unexpected data may result in L1 cache corruption
P25	X						Fixed	Multi-processor boot protocol may not complete with an IOQ depth of one
P26	X	X	X	X	X	X	No Fix	Processor flags #PF instead of #AC on an unlocked CMPXCHG8B instruction
P27	X	X	X				Fixed	Incorrect data may be returned when page tables are located in Write Combining (WC) memory
P28	X	X	X	X	X	X	No Fix	FSW may not be completely restored after page fault on FRSTOR or FLDENV instructions
P29	X	X	X				Fixed	Write Combining (WC) load may result in bogus address on system bus
P30	X	X	X	X	X	X	No Fix	Processor provides a 4-byte store unlock after an 8-byte load lock
P31	X	X	X				Fixed	Multiple accesses to the same S-state L2 cache line and ECC error combination may result in loss of cache coherency
P32	X	X	X	X	X	X	No Fix	IA32_MC0_ADDR and IA32_MC0_MISC registers will contain invalid or stale data following a data, address, or response parity error
P33	X	X	X	X	X	X	No Fix	When the processor is in the System Management Mode (SMM), Debug Registers may be fully writeable
P34	X	X	X	X	X	X	No Fix	Associated counting logic must be configured when using Event Selection Control (ESCR) MSR
P35	X	X	X				Fixed	Livelock may occur when bus parking is disabled
P36	X	X	X				Fixed	CR2 may be incorrect or an incorrect page fault error code may be pushed onto stack after execution of an LSS instruction
P37	X	X					Fixed	Buffer On Resistance may exceed specification
P38			X				Fixed	Instruction pointer stored on stack may become invalid
P39			X	X	X	X	No Fix	Shutdown and IERR# may result due to a Machine Check Exception on a Hyper-Threading technology enabled processor
P40			X				Fixed	Hyper-Threading technology enabled processors may hang in the presence of extensive self-modifying code
P41			X				Fixed	Global bit incorrectly set for secondary logical processors in ITLB
P42			X				Fixed	Machine Check Exception (MCE) observed on DP Platforms
P43			X	X	X	X	Plan Fix	BPM[5:3]# V _{IL} does not meet specification
P44			X	X	X	X	No Fix	Processor may hang under certain frequencies and 12.5% STPCLK# duty cycle

SUMMARY OF ERRATA								
NO.	C1/ 0F0Ah	D0/ 0F12h	B0/ 0F24h	C1/ 0F27h	D1/ 0F29h	M0/ 0F25h	PLANS	ERRATA
P45	X	X	X	X	X	X	No Fix	System may hang if a fatal cache error causes Bus Write Line (BWL) transaction to occur to the same cache line address as an outstanding Bus Read Line (BRL) or Bus Read-Invalidate Line (BRIL)
P46	X	X	X				Fixed	L2 Cache may contain Stale Data in the Exclusive State
P47	X	X	X	X	X	X	No Fix	Re-mapping the APIC Base Address to a value less than or equal to 0xDC001000 may cause IO and Special Cycle Failure
P48			X	X	X	X	Plan Fix	Erroneous BIST result found in EAX Register after Reset
P49	X	X	X				Fixed	Processor does not Flag #GP on no-zero write to certain MSRs
P50	X	X	X	X	X		No Fix	Simultaneous assertion of A20M# and INIT# may result in incorrect data fetch
P51			X	X			Fixed	Processor does not respond to Break Requests from ITP
P52	X	X	X				Fixed	CPUID instruction returns incorrect number of ITLB entries
P53	X	X	X	X	X	X	No Fix	A write to APIC Task Priority Register (TPR) that lowers priority may seem to have not occurred
P54	X	X					Fixed	Stop-Clock Assertion May Cause a System to Hang
P55				X	X	X	Plan Fix	Store to Load Data Forwarding may result in switched data bytes
P56				X	X	X	Plan Fix	ITP cannot continue Single Step Execution after the first breakpoint
P57	X	X	X	X	X	X	No Fix	Parity Error in the L1 Cache may cause the processor to hang
P58	X	X	X	X	X	X	Plan Fix	The TCK Input in the Test Access Port (TAP) is Sensitive to Low Clock Edge-rates and Prone to Noise Coupling Onto TCK's Rising or Falling Edges
P59	X	X	X	X	X	X	No Fix	Disabling a Local APIC Disables Both Logical Processor APICs on a Hyper-Threading Technology Enabled Processor
P60	X	X	X	X	X	X	No fix	Using STPCLK and Executing Code From Very Slow Memory Could Lead to a System Hang
P61						X	No Fix	Simultaneous cacheline eviction from the L2 and the L3 caches may result in the write back of stale data.
P62	X	X	X	X	X	X	No Fix	The State of the Resume Flag (RF Flag) in a Task-State Segment (TSS) May be Incorrect
P63		X		X	X		No Fix	Changes to CR3 Register do not Fence Pending Instruction Page Walks
P64			X	X	X	X	Plan Fix	Simultaneous Page Faults at Similar Page Offsets on Both Logical Processors of an Hyper-Threading Technology Enabled Processor May Cause Application Failure



Number	Plan	SPECIFICATION CHANGES
P1	Doc	Context ID feature added to Processor Signature /IA32_MISC_ENABLE registers

Number	Plan	SPECIFICATION CLARIFICATIONS
P1	Doc	Thermal Specification Clarification – Intel Xeon Processor
P2	Doc	Thermal Specification Clarification – Intel Xeon Processor with 512-KB L2 cache
P4	Doc	Thermal Specification Clarification - Low Voltage Intel Xeon Processor

Errata

P1. UC Code in Same Line as WriteBack (WB) Data may Lead to Data Corruption

Problem: This erratum occurs when both code (being accessed as UC or WC) and data (being accessed as WB) are placed in the same cache line. The UC fetch will cause the processor to self-snoop and generate an implicit writeback. The data supplied by this implicit writeback may be corrupted due to the way the processor is currently handling self-modifying code.

Implication: UC code located in the same cache line as WB data may lead to data corruption.

Workaround: UC or WC code should not be located in the same 64 byte cache line as any location that is being stored to with WB data.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P2. Transaction Is Not Retried After BINIT#

Problem: If the first transaction of a locked sequence receives a HITM# and DEFER# during the snoop phase it should be retried and the locked sequence restarted. However, if BINIT# is also asserted during this transaction, the transaction will not be retried.

Implication: When this erratum occurs, locked transactions will not be retried.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P3. Invalid Opcode 0FFFh Requires A ModRM Byte

Problem: Some invalid opcodes require a ModRM byte and other following bytes, while others do not. The invalid opcode 0FFFh did not require a ModRM in previous generation microprocessors such as Pentium® II or Pentium III processors, but it is required in the Intel® Xeon™ processor

Implication:

Implication: The use of an invalid opcode 0FFFh without the ModRM byte may result in a page or limit fault on the

Implication: Intel Xeon processor.

Workaround: To avoid this erratum use ModRM byte with invalid 0FFFh opcode.

Status: For the steppings affected, see the *Summary Tables of Changes*.



P4. When in No-Fill Mode (CR0.CD=1) the Memory Type of Large (PSE-4M and PAE-2M) Pages are Wrongly Forced to Uncacheable

Problem: When the processor is operating in No-Fill Mode (CR0.CD=1), the page miss hardware incorrectly forces the memory type of large (PSE-4M and PAE-2M) pages to UC memory type regardless of the MTRR settings. By forcing the memory type of these pages to UC, load operations, which should hit valid data in the L1 cache, are forced to load the data from system memory. Some applications will lose the performance advantage associated with the caching permitted by other memory types.

Implication: This erratum may result in some performance degradation when using no-fill mode with large pages.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P5. Processor May Hang Due to Speculative Page Walks to NonExistent System Memory

Problem: A load operation issued speculatively by the processor that misses the Data Translation Lookaside Buffer (DTLB) results in a page walk. A branch instruction older than the load retires so that this load operation is now in the mispredicted branch path. Due to an internal boundary condition, in some instances the load is not canceled before the page walk is issued.

The Page Miss Handler (PMH) starts a speculative page-walk for the Load and issues a cacheable load of the Page Directory Entry (PDE). This PDE load returns data that points to a page table entry in uncacheable (UC) memory. The PMH issues the PTE Load to UC space, which is issued on the Front Side Bus. No response comes back for this load PTE operation since the address is pointing to system memory, which does not exist.

This load to non-existent system memory causes the processor to hang because other bus requests are queued up behind this UC PTE load, which never gets a response. If the load was accessing valid system memory, the speculative page-walk would successfully complete and the processor would continue to make forward progress.

Implication: Processor may hang due to speculative page walks to non-existent system memory

Workaround: Page directories and page tables in UC memory space must point to system memory that exists.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P6. Writing a Performance Counter May Result in an Incorrect Counter Value

Problem: Accessing a performance counter also enables the counter input so that writing one half of the counter can cause the other half to increment. When a performance counter is written and the event counter for the event being monitored is non-zero, the performance counter will be incremented by the value on that event counter. Because the upper eight bits of the performance counter are not written at the same time as the lower 32 bits, the increment due to the non-zero event counter may cause a carry to the upper bits such that the performance counter contains a value higher than what was written. The worst-case error caused by this can be about 4 billion counts.

Implication: When this erratum occurs, the performance counter will contain a different value from that which was written.

Workaround: If the performance counter is set to select a null event and the CCCR for that counter has its compare bit set to zero, before the performance counter is written, this erratum will not occur.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P7. Performance Counter May Contain Incorrect Value After Being Stopped

Problem: If a performance counter is stopped on the precise internal clock cycle where the intermediate carry from the lower 32 bits of the counter to the upper eight bits occurs, the intermediate carry is lost.

Implication: When this erratum occurs the performance counter may contain a value about 4 billion (2³²) less than it should.

Workaround: Since this erratum does not occur if the performance counters are read when running, a possible workaround is to read the counter before stopping it.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P8. REP MOV Instruction with Overlapping Source and Destination may Result in Data Corruption

Problem: When fast strings are enabled and a REP MOV instruction is used to move a string and the source and destination strings overlap by 56 bytes or less, data corruption may occur.

Implication: When this erratum occurs, data corruption may occur.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P9. Memory Type of the Load Lock Different from its Corresponding Store Unlock

Problem: A use-once protocol is employed to ensure that the processor in a multi-agent system may access data that is loaded into its cache on a Read-for-Ownership operation at least once before it is snooped out by another agent. This protocol is necessary to avoid a multi-agent livelock scenario in which the processor cannot gain ownership of a line and modify it before that data is snooped out by another agent. In the case of this erratum, split load lock instructions incorrectly trigger the use-once protocol. A load lock operation accesses data that splits across a page boundary with both pages of WB memory type. The use-once protocol activates and the memory type for the split halves get forced to UC. Since use-once does not apply to stores, the store unlock instructions go out as WB memory type. The full sequence on the bus is: locked partial read (UC), partial read (UC), partial write (WB), locked partial write (WB). The use-once protocol should not be applied to load locks.

Implication: When this erratum occurs, the memory type of the load lock will be different than the memory type of the store unlock operation. This behavior (load locks and store unlocks having different memory types) does not introduce any functional failures such as system hangs or memory corruption.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P10. Machine Check Architecture Error Reporting and Recovery May Not Work as Expected

Problem: When the processor detects errors it should attempt to report and/or recover from the error. In the situations described below, the processor does not report and/or recover from the error(s) as intended.

When a transaction is deferred during the snoop phase and subsequently receives a Hard Failure response, the transaction should be removed from the bus queue so that the processor may proceed. Instead, the transaction is not properly removed from the bus queue, the bus queue is blocked, and the processor will hang.

When a hardware prefetch results in an uncorrectable tag error in the L2 cache, MC0_STATUS.UNCOR and MC0_STATUS.PCC are set but no Machine Check Exception (MCE) is signaled. No data loss or corruption occurs because the data being prefetched has not been used. If the data location with the uncorrectable tag error is subsequently accessed, an MCE will occur. However, upon this MCE, or any other subsequent MCE, the information for that error will not be logged because MC0_STATUS.UNCOR has already been set and the MCA status registers will not contain information about the error which caused the MCE assertion but instead will contain information about the prefetch error event.

When the reporting of errors is disabled for Machine Check Architecture (MCA) Bank 2 by setting all MC2_CTL register bits to 0, uncorrectable errors should be logged in the IA32_MC2_STATUS register but no machine-check exception should be generated. Uncorrectable loads on bank 2, which would normally be logged in the IA32_MC2_STATUS register, are not logged.

When one half of a 64 byte instruction fetch from the L2 cache has an uncorrectable error and the other 32 byte half of the same fetch from the L2 cache has a correctable error, the processor will attempt to correct the correctable error but cannot proceed due to the uncorrectable error. When this occurs the processor will hang.

When an L1 cache parity error occurs, the cache controller logic should write the physical address of the data memory location that produced that error into the IA32_MC1_ADDR REGISTER (MC1_ADDR). In some instances of a parity error on a load operation that hits the L1 cache, however, the cache controller logic may write the physical address from a subsequent load or store operation into the IA32_MC1_ADDR register.

The local xAPIC has an Error Status Register, which records all errors it detects. Bit 6 of this register, the Receive Illegal Vector bit, is set when the local xAPIC detects an illegal vector in a message that it received. When an illegal vector error is received on the same internal clock that the error status register is being written due to a previous error, bit 6 does not get set and illegal vector errors are not flagged.

When an error exists in the tag field of a cache line such that a request for ownership (RFO) issued by the processor hits multiple tag fields in the L2 cache (the correct tag and the tag with the error) and the accessed data also has a correctable error, the processor will correctly log the multiple tag match error but will hang when attempting to execute the machine check exception handler.

If a memory access receives a machine check error on both 64 byte halves of a 128-byte L2 cache sector, the IA32_MC0_STATUS register records this event as multiple errors, i.e., the valid error bit and the overflow error bit are both set indicating that a machine check error occurred while the results of a previous error were in the error-reporting bank. The IA32_MC1_STATUS register should also record this event as multiple errors but instead records this event as only one correctable error.

The overflow bit should be set to indicate when more than one error has occurred. The overflow bit being set indicates that more than one error has occurred. Because of this erratum, if any further errors occur, the MCA overflow bit will not be updated, thereby incorrectly indicating only one error has been received.

If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the processor will signal a Machine Check Exception (MCE). If the instruction is directed at a device that is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, while attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is successfully completed, it will attempt to restart the I/O instruction, but will not have the correct machine state due to the call to the MCE handler. This can lead to failure of the restart and shutdown of the processor.

If PWRGOOD is de-asserted during a RESET# assertion causing internal glitches, the MCA registers may latch invalid information.

If RESET# is asserted, then de-asserted, and reasserted, before the processor has cleared the MCA registers, then the information in the MCA registers may not be reliable, regardless of the state or state transitions of PWRGOOD.

If MCERR# is asserted by one processor and observed by another processor, the observing processor does not log the assertion of MCERR#. The Machine Check Exception (MCE) handler called upon assertion of MCERR# will not have any way to determine the cause of the MCE.

The Overflow Error bit (bit 62) in the IA32_MC0_STATUS register indicates, when set, that a machine check error occurred while the results of a previous error were still in the error reporting bank (i.e. The Valid bit was set when the new error occurred). If an uncorrectable error is logged in the error-reporting bank and another error occurs, the overflow bit will not be set.

The MCA Error Code field of the IA32_MC0_STATUS register gets written by a different mechanism than the rest of the register. For uncorrectable errors, the other fields in the IA32_MC0_STATUS register are only updated by the first error. Any further errors that are detected will update the MCA Error Code field without updating the rest of the register, thereby leaving the IA32_MC0_STATUS register with stale information.

When a speculative load operation hits the L2 cache and receives a correctable error, the IA32_MC1_Status Register may be updated with incorrect information. The IA32_MC1_Status Register should not be updated for speculative loads.

The processor should only log the address for L1 parity errors in the IA32_MC1_Status register if a valid address is available. If a valid address is not available, the Address Valid bit in the IA32_MC1_Status register should not be set. In instances where an L1 parity error occurs and the address is not available because the linear to physical address translation is not complete or an internal resource conflict has occurred, the Address Valid bit is incorrectly set.

The processor may hang when an instruction code fetch receives a hard failure response from the Front Side Bus. This occurs because the bus control logic does not return data to the core, leaving the processor empty. IA32_MC0_STATUS MSR does indicate that a hard fail response occurred.

The processor may hang when the following events occur and the machine check exception is enabled, CR4.MCE=1. A processor that has its STPCLK# pin asserted will internally enter the Stop Grant State and finally issue a Stop Grant Acknowledge special cycle to the bus. If an uncorrectable error is generated during the Stop Grant process it is possible for the Stop Grant special cycle to be issued to the bus before the processor vectors to the machine check handler. Once the chipset receives its last Stop Grant special cycle it is allowed to ignore any bus activity from the processors. As a result, processor accesses to the machine check handler may not be acknowledged, resulting in a processor hang.

Implication: The processor is unable to correctly report and/or recover from certain errors

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P11. Debug Mechanisms May Not Function as Expected

Problem: Certain debug mechanisms may not function as expected on the processor. The cases are as follows:

When the following conditions occur: 1) An FLD instruction signals a stack overflow or underflow, 2) the FLD instruction splits a page-boundary or a 64 byte cache line boundary, 3) the instruction matches a Debug Register on the high page or cache line respectively, and 4) the FLD has a stack fault and a memory fault on a split access, the processor will only signal the stack fault and the debug exception will not be taken.

When a data breakpoint is set on the ninth and/or tenth byte(s) of a floating point store using the Extended Real data type, and an unmasked floating point exception occurs on the store, the breakpoint will not be captured.

When any instruction has multiple debug register matches, and any one of those debug registers is enabled in DR7, all of the matches should be reported in DR6 when the processor goes to the debug handler. This is not true during a REP instruction. As an example, during execution of a REP MOVSW instruction the first iteration a load matches DR0 and DR2 and sets DR6 as FFFF0FF5h. On a subsequent iteration of the instruction, a load matches only DR0. The DR6 register is expected to still contain FFFF0FF5h, but the processor will update DR6 to FFFF0FF1h.

A Data breakpoint that is set on a load to uncacheable memory may be ignored due to an internal segment register access conflict. In this case the system will continue to execute instructions, bypassing the intended breakpoint. Avoiding having instructions that access segment descriptor registers e.g. LGDT, LIDT close to the UC load, and avoiding serialized instructions before the UC load will reduce the occurrence of this erratum.

Implication: Certain debug mechanisms do not function as expected on the processor.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P12. Processor May Live-lock if PDEs or PTEs are in UC Space

Problem: The processor may livelock under the following boundary conditions:

The Page-Directory Entries (PDEs) or Page-Table Entries (PTEs) are in uncacheable (UC) space

An instruction fetch misses the ITLB resulting in a page walk

This instruction fetch is immediately followed by a store that splits a page boundary

Implication: When this erratum occurs, the processor will livelock. This erratum was found using random instruction testing and has not been observed with commercial software.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P13. Thermal Status Log Bit May Not be Set when the Thermal Control Circuit is Active

Problem: Bit 1 of the IA32_THERM_STATUS register (Thermal Status Log) is a sticky bit designed to be set to '1' if the thermal control circuit (TCC) has been active since either the previous processor reset or software cleared this bit. If TCC is active and the Thermal Status Log bit is cleared by a processor reset or by software, it will remain clear (set to '0') as long as the TCC remains active. Once TCC deactivates, the next activation of the TCC will set the Thermal Status Log bit.

Implication: When this erratum occurs, the Thermal Status Log bit will be cleared (set to '0') although the thermal control circuit is active.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*

P14. Processor May Timeout Waiting For A Device To Respond After 0.67 Seconds

Problem: The PCI 2.1 target initial latency specification allows two seconds for a device to respond during initialization-time. The processor may timeout after only approximately 0.67 seconds. When the processor times out it will hang with IERR# asserted. PCI devices that take longer than 0.67 seconds to initialize may not be initialized properly.

Implication: System may hang with IERR# asserted.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

Status:

P15. Cascading Of Performance Counters Does Not Work Correctly When Forced Overflow is Enabled

Problem: The performance counters are organized into pairs. When the CASCADE bit of the Counter Configuration Control Register (CCCR) is set, a counter that overflows will continue to count in the other counter of the pair. The FORCE_OVF bit forces the counters to overflow on every non-zero increment. When the FORCE_OVF bit is set, the counter overflow bit will be set but the counter no longer cascades.

Implication: The performance counters do not cascade when the FORCE_OVF bit is set.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P16. EMON Event Counting of x87 Loads May Not Work as Expected

Problem: If a performance counter is set to count x87 loads and floating-point exceptions are unmasked, the FPU Operand (Data) Pointer (FDP) may become corrupted.

Implication: When this erratum occurs, FPU Operand (Data) Pointer (FDP) may become corrupted.

Workaround: This erratum will not occur with floating point exceptions masked. If floating-point exceptions are unmasked, then performance counting of x87 loads should be disabled.

Status: For the steppings affected, see the *Summary Tables of Changes*.



P17. Simultaneous Code Breakpoint and Uncorrectable Error Results in Processor Hang

Problem: If an instruction fetch results in an uncorrectable error and there is also a debug breakpoint at this address, the processor will hang and the uncorrectable error will not be logged in the Machine Check registers.

Implication: When this erratum occurs the processor will hang.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P18. Software Controlled Clock Modulation Using A 12.5% or 25% Duty Cycle May Cause The Processor To Hang

Problem: Per the ACPI 1.0b specification, processor clock modulation may be controlled via a processor register (IA32_THERM_CONTROL). The On-Demand Clock Modulation Duty Cycle is controlled by bits 3:1. If these bits are set to a duty cycle of 12.5% or 25%, the processor may hang while attempting to execute a floating-point instruction. In this failure, the last instruction pointer (LIP) is pointing to a floating-point instruction whose instruction bytes are in UC space and which takes an exception 16 (floating point error exception). The processor stalls trying to fetch the bytes of the faulting floating-point instruction and those following it. This processor hang is caused by interactions between the thermal control circuit and floating-point event handler.

Implication: When the clock modulation is set to 12.5% or 25% duty cycle, the processor will go into a sleep state from which it fails to return.

Workaround: Use a duty cycle other than 12.5% or 25%.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P19. RFO With ECC Error May Result In Incorrect Data

Problem: This erratum occurs as the result of the following conditions:

A read for ownership (RFO) generates a correctable error.

In the process of correcting the error, a locked RFO (LRFO) is issued that uses the same internal buffer as the previous RFO.

Another processor issues a snoop to the same address as the LRFO.

An internal boundary condition exists which may prevent the LRFO from completing correctly causing the snoop to receive incorrect data. Intel has not been able to reproduce this erratum with commercial software.

Implication: When this erratum occurs, data corruption may result.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P20. Speculative Page Fault May Cause Livelock

Problem: If the processor detects a page fault, which is corrected before the operating system page fault handler can be called (e.g. a second processor or DMA activity modifies the page tables and the corrected page tables are left in a non-accessed or non-modified state) the processor may livelock. Intel has not been able to reproduce this erratum with commercial software.

Implication: This erratum occurs in systems where page tables are being modified by other processors. If this erratum is encountered, the processor will livelock resulting in a system hang or operating system failure.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P21. PAT Index MSB May Be Calculated Incorrectly

Problem: When Mode B or Mode C paging support is enabled and all of the following events occur:

A page walk returns the Page Directory Entry (PDE) for a large page from memory.

A subsequent page walk returns the Page Table Entry (PTE) for a 4k page from memory and the Page Attribute Table (PAT) upper index bit in this PTE is set to 1b.

It is possible that the PAT upper index bit in the PTE is incorrectly ignored and assumed to be 0b. The result is that the memory type in the PAT that should have come from the corresponding PAT index [4-7] incorrectly comes from PAT index [0-3].

Implication: If an operating system has programmed the PAT in an asymmetrical fashion i.e. PAT[0-3] is different from PAT[4-7] then an incorrect memory type may be used.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.

P22. System Bus Interrupt Messages Without Data And Which Receive A HardFailure Response May Hang The Processor

Problem: When a system bus agent (processor or chipset) issues an interrupt transaction without data onto the system bus and the transaction receives a HardFailure response, a potential processor hang can occur. The processor, which generates an inter-processor interrupt (IPI) that receives the HardFailure response, will still log the MCA error event cause as HardFailure, even if the APIC causes a hang. Other processors, which are true targets of the IPI, will also hang on hardfail-without-data, but will not record an MCA HardFailure event as the cause. If a HardFailure response occurs on a system bus interrupt message with data, the APIC will complete the operation so as not to hang the processor.

Implication: The processor may hang.

Workaround: None identified

Status: For the steppings affected, see the *Summary Tables of Changes*.



P23. SQRTPD and SQRTPD May Return QNaN Indefinite Instead of Negative Zero

Problem: When DAZ mode is enabled, and a SQRTPD or SQRTPD instruction has a negative denormal operand, the instruction will return a QNaN indefinite when the specified response should be zero.

Implication: When this erratum occurs, the instruction will return a QNaN indefinite when a zero is expected.

Workaround: Ensure that negative denormals are not used as operands to the SQRTPD or SQRTPD instructions when DAZ mode is enabled. Software could enable FTZ mode to ensure that negative denormals are not generated by computation prior to execution of the SQRTPD or SQRTPD instructions.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P24. Bus Invalidate Line Requests that Returns Unexpected Data May Result in L1 Cache Corruption

Problem: When a Bus Invalidate Line (BIL) request receives unexpected data from a deferred reply, and a store operation write combines to the same address, there is a small window where the L1 cache is corrupt, and loads can retire with this corrupted data. This erratum occurs in the following scenario:

A Read-For-Ownership (RFO) transaction is issued by the processor and hits a line in shared state in the L2 cache.

The RFO is then issued on the System Bus as a 0 length Read-Invalidate (BIL), since it doesn't need data, just ownership of the cache line.

This transaction is deferred by the chipset.

At some later point, the chipset sends a deferred reply for this transaction with an implicit write-back response. For this erratum to occur, no snoop of this cache line can be issued between the BIL and the deferred reply.

The processor issues a write-combining store to the same cache line while data is returning to the processor. This store straddles an 8-byte boundary.

Note: Due to an internal boundary condition, a time window exists where the L1 cache contains corrupt data, which could be accessed by a load.

Implication: The L1 cache may contain corrupted data. No known commercially available chipsets trigger the failure conditions.

Workaround: The chipset could issue a BIL (snoop) to the deferred processor to eliminate the failure conditions.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P25. Multi-processor Boot Protocol may not Complete with an IOQ Depth of One

Problem: When the In-Order Queue (IOQ) depth is managed by the chipset to be one entry deep, the system may hang during the multi-processor boot protocol. This hang occurs when the chipset drives BNR# in such a way that the processors are continually throttled off the bus then released to access the bus in alternating cycles which never allows the multi-processor boot protocol to complete execution.

Implication: The system may hang during the multi-processor boot protocol.

Workaround: If the chipset drives BNR# in such a way that the processors are continually throttled off the bus then released to access the bus in alternating cycles, do not use In-Order Queue de-pipelining.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P26. Processor Flags #PF Instead of #AC on an Unlocked CMPXC8B Instruction

Problem: If a data page fault (#PF) and alignment check fault (#AC) both occur for an unlocked CMPXC8B instruction, then #PF will be flagged.

Implication: Software that depends #AC before #PF will be affected since #PF is flagged in this case.

Workaround: Remove the software's dependency on the fact that #AC has precedence over #PF. Alternately, if the reload is due to a not present page, reload the page in the page fault handler and then restart the faulting instruction.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P27. Incorrect Data may be Returned When Page Tables are Located in Write Combining (WC) Memory

Problem: If page directories and/or page tables are located in Write Combining (WC) memory, speculative loads to cacheable memory may complete with incorrect data.

Implication: Cacheable loads to memory mapped using page tables located in write combining memory may return incorrect data. Intel has not been able to reproduce this erratum with commercially available software.

Workaround: Do not place page directories and/or page tables in WC memory.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P28. FSW May not be Completely Restored after Page Fault on FRSTOR or FLDENV Instructions

Problem: If the FPU operating environment or FPU state (operating environment and register stack) being loaded by an FLDENV or FRSTOR instruction wraps around a 64-Kbyte or 4-Gbyte boundary and a page fault (#PF) or segment limit fault (#GP or #SS) occurs on the instruction near the wrap boundary, the upper byte of the FPU status word (FSW) might not be restored. If the fault handler does not restart program execution at the faulting instruction, stale data may exist in the FSW.

Implication: When this erratum occurs, stale data will exist in the FSW.

Workaround: Ensure that the FPU operating environment and FPU state do not cross 64-Kbyte or 4-Gbyte boundaries. Alternately, ensure that the page fault handler restarts program execution at the faulting instruction after correcting the paging problem.

Status: For the steppings affected, see the *Summary Tables of Changes*.



P29. Write Combining (WC) Load May Result in an Unintended Address on System Bus

Problem: When the processor performs a speculative write combining (WC) load, down the path of a mispredicted branch, and the address happens to match a valid UnCacheable (UC) address translation with the Data Translation Look-Aside Buffer, an unintended UnCacheable load operation may be sent out on the System Bus.

Implication: When this erratum occurs, an unintended load may be sent on System Bus. Intel has only encountered this erratum during pre-silicon simulation.

Workaround: It is possible for the BIOS to contain a workaround for this erratum for some steppings of the processor.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P30. Processor provides a 4-byte store unlock after an 8-byte load lock

Problem: When the processor is in the Page Address Extension (PAE) mode and detects the need to set the Access and/or Dirty bits in the page directory or page table entries, the processor sends an 8 byte load lock onto the System Bus. A subsequent 8 byte store unlock is expected, but instead a 4 byte store unlock occurs. Correct data is provided since only the lower bytes change, however external logic monitoring the data transfer may be expecting an 8-byte store unlock.

Implication: No known commercially available chipsets are affected by this erratum.

Workaround: None identified at this time.

Status: For the steppings affected, see the *Summary Tables of Changes*

Status: .

P31. Multiple Accesses to the Same S-State L2 Cache Line and ECC Error Combination May Result in Loss of Cache Coherency

Problem: When a Read For Ownership (RFO) cycle has a 64 bit address match with an outstanding read hit on a line in the L2 cache which is in the S-state AND that line contains an ECC error, the processor should recycle the RFO until the ECC error is handled. Due to this erratum, the processor does not recycle the RFO and attempt to service both the RFO and the read hit at the same time.

Implication: When this erratum occurs, cache may become incoherent.

Workaround: None identified at this time.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P32. IA32_MC0_ADDR and IA32_MC0_MISC Registers Will Contain Invalid or Stale Data Following a Data, Address, or Response Parity Error

Problem: If the processor experiences a data, address, or response parity error, the ADDR_V and MISC_V bits of the IA32_MC0_STATUS register are set, but the IA32_MC0_ADDR and IA32_MC0_MISC registers are not loaded with data regarding the error.

Implication: When this erratum occurs, the IA32_MC0_ADDR and IA32_MC0_MISC registers will contain invalid or stale data.

Workaround: Ignore any information in the IA32_MC0_ADDR and IA32_MC0_MISC registers after a data or response parity error.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P33. When the Processor is in the System Management Mode (SMM), Debug Registers may be Fully Writeable

Problem: When in System Management Mode (SMM), the processor executes code and stores data in the SMRAM space. When the processor is in this mode and writes are made to DR6 and DR7, the processor should block writes to the reserved bit locations. Due to this erratum, the processor may not block these writes. This may result in invalid data in the reserved bit locations.

Implication: Reserved bit locations within DR6 and DR7 may become invalid.

Workaround: Software may perform a read/modify/write when writing to DR6 and DR7 to ensure that the value in the reserved bits are maintained.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P34. Associated Counting Logic must be Configured when Using Event Selection Control (ESCR) MSR

Problem: ESCR MSRs allow software to select specific events to be counted, with each ESCR usually associated with a pair of performance counters. ESCRs may also be used to qualify the detection of at-retirement events that support precise-event-based sampling (PEBS). A number of performance metrics that support PEBS require a 2nd ESCR to tag uops for the qualification of at-retirement events. (The first ESCR is required to program the at-retirement event.) Counting is enabled via counter configuration control registers (CCCR) while the event count is read from one of the associated counters. When counting logic is configured for the subset of at-retirement events that require a 2nd ESCR to tag uops, at least one of the CCCRs in the same group of the 2nd ESCR must be enabled.

Implication: If no CCCR/counter is enabled in a given group, the ESCR in that group that is programmed for tagging uops will have no effect. Hence a subset of performance metrics that require a 2nd ESCR for tagging uops may result in 0 count.

Workaround: Ensure that at least one CCCR/counter in the same group as the tagging ESCR is enabled for those performance metrics that require 2 ESCRs and tagging uops for at-retirement counting.

Status: For the steppings affected, see the *Summary Tables of Changes*.



P35. Livelock May Occur When Bus Parking Is Disabled

Problem: A livelock may occur when processor bus parking is disabled, and when (1) the processor is the symmetric owner of the bus with one internal request pending, and (2) the processor observes the assertion of BPRI#, BNR# or a full In Order Queue (IOQ). In this scenario, the processor bus interface unit assumes that the assertion of ADS# is not required, deasserts BREQ, and, as a result, relinquishes bus ownership without issuing the pending request. If the BPRI#, BNR# or full IOQ pattern continues coincident with the arbitration phase of the processor that still has only one outstanding internal request, livelock may occur. Assertion of bus parking, any change to the regular pattern of BPRI# or BNR# assertion noted above, or the arrival of a second internal transaction will release the processor from the livelock condition.

Implication: This erratum may result in a livelock.

Workaround: This erratum can be avoided by enabling bus parking. The deassertion of signal A15# during the active-to-inactive edge of RESET# will enable bus parking.

Status: For the stepping affected, see the *Summary of Changes* at the beginning of this section.

Status:

P36. CR2 May Be Incorrect or an Incorrect Page Fault Error Code May Be Pushed onto Stack After Execution of an LSS Instruction

Problem: Under certain timing conditions, the internal load of the selector portion of the LSS instruction may complete with potentially incorrect speculative data before the load of the offset portion of the address completes. The incorrect data is corrected before the completion of the LSS instruction but the value of CR2 and the error code pushed on the stack are reflective of the speculative state. Intel has not observed this erratum with commercially available software.

Implication: When this erratum occurs, the contents of CR2 may be off by two, or an incorrect page fault error code may be pushed onto the stack.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P37. Buffer on Resistance May Exceed Specification

Problem: The datasheet specifies the resistance range for RON (Buffer On Resistance) for the AGTL+ buffer as 5 to 11 ohms. Due to this erratum, RON may be as high as 12 ohms.

Implication: The RON value affects the voltage level of the signals when the buffer is driving the signal low. A higher RON may adversely affect the system's ability to meet specifications such as VIL. As the system design also affects margin to specification, designs may or may not have sufficient margin to function properly with an increased RON. System designers should evaluate whether a particular system is affected by this erratum. Designs that follow the recommendations in the Intel® Xeon™ Processor and Intel® 860 Chipset Platform Design Guide are not expected to be affected.

Workaround: No workaround is necessary for systems with margin sufficient to accept a higher R_{ON}.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P38. Instruction Pointer Stored on Stack May Become Invalid

Problem: The instruction pointer stored on the stack may become invalid due to an internal boundary condition which may exist on a Hyper Threading Technology enabled processors. The following sequence of events must occur in order to encounter this erratum:

1. One logical processor executes the WRMSR instruction with incorrect data causing a general protection fault
2. Simultaneously, an event that requires micro-architectural synchronization among the two logical processors occurs on the second logical processor. This event may cause an invalid instruction pointer to be stored on the ring 0 stack during the transition to GP fault handler on the first logical processor.

Implication: The instruction pointer stored on the stack may be invalid, potentially causing errors during execution of or return from the GP fault handler.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P39. Shutdown and IERR# May Result Due to a Machine Check Exception on a Hyper-Threading Technology Enabled Processor

Problem: When a Machine Check Exception (MCE) occurs due to an internal error, both logical processors on a Hyper-Threading Technology enabled processor normally vector to the MCE handler. However, if one of the logical processors is in the “Wait-for-SIPI” state, that logical processor will not have an MCE handler and will shut down and assert IERR#.

Implication: A processor with a logical processor in the “Wait for SIPI” state will shut down when an MCE occurs on the other thread.

Workaround: None identified at this time.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P40. Hyper-Threading Technology Enabled Processors May Hang in the Presence of Extensive Self-Modifying Code

Problem: For multi-processor platforms, in which Hyper-Threading Technology enabled processors are executing extensive self modifying code, and branch trace messages are enabled on at least one logical processor, the system may hang. In this scenario, a processor executing within 1K of code being written to by another processor may attempt to end this flow, thereby resulting in a hang.

Implication: When this erratum occurs the system will hang.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.



P41. Global Bit Incorrectly Set for Secondary Logical Processors in ITLB

Problem: Due to a boundary condition in the translation look-aside buffer logic, the global bit information in the TLB entry for a mapping belonging to the first logical processor can overwrite the global bit information for a mapping belonging to the second logical processor. This occurs in the following scenario:

The first logical processor misses the ITLB resulting in a page walk

The second logical processor also misses the ITLB and generates a page walk.

Note: In certain timing scenarios within the processor, the leftover global bit information from the first logical processor may overwrite the second logical processor.

Implication: When this erratum occurs, if the Page global bit for the second logical processor is overwritten with a 0b, this will result in performance degradation for the first logical processor. If the page global bit is incorrectly changed from a 0 to 1, this erratum may result in software failures.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P42. Machine Check Exception (MCE) Observed on DP Platforms

Problem: A System Bus address parity error may be signaled if two processors run at odd core frequency to System Bus-ratios (17:1, 19:1, etc) on DP processor platforms. This address parity error signaling issue does not occur if the processors run at even bus-ratios.

Implication: A Machine Check Exception (MCE) may be observed on DP platforms.

Workaround: The system BIOS should ensure that processors run at even core frequency to System Bus-ratios (16:1, 18:1, etc).

Status: For the steppings affected, see the *Summary Tables of Changes*.

P43. BPM[5:3]# VIL Does Not Meet Specification

Problem: The VIL for BPM[5:3]# is specified as $0.9 * GTLREF$ [V]. Due to this erratum the VIL for these signals is $0.9 * GTLREF - .100$ [V].

Implication:

Implication: The processor requires a lower input voltage than specified to recognize a low voltage on the BPM[5:3]#

Implication: signals.

Workaround: When intending to drive the BPM[5:3]# signals low, ensure that the system provides a voltage lower than $0.9 * GTLREF - .100$ [V].

Status: For the steppings affected, see the *Summary Tables of Changes*.

P44. Processor May Hang Under Certain Frequencies and 12.5% STPCLK# Duty Cycle

Problem: If a system de-asserts STPCLK# at a 12.5% duty cycle, and the processor is running below 2 GHz, and the processor thermal control circuit (TCC) on-demand clock modulation is active, the processor may hang. This erratum does not occur under the automatic mode of the TCC.

Implication: When this erratum occurs, the processor will hang.

Workaround: If use of the on-demand mode of the processor's TCC is desired in conjunction with STPCLK# modulation, then assure that STPCLK# is not asserted at a 12.5% duty cycle.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P45. System May Hang if a Fatal Cache Error Causes Bus Write Line (BWL) Transaction to Occur to the Same Cache Line Address as an Outstanding Bus Read Line (BRL) or Bus Read-Invalidate Line (BRIL)

Problem: A processor internal cache fatal data ECC error may cause the processor to issue a BWL transaction to the same cache line address as an outstanding BRL or BRIL. As it is not typical behavior for a single processor to have a BWL and a BRL/BRIL concurrently outstanding to the same address, this may represent an unexpected scenario to system logic within the chipset.

Implication: The processor may not be able to fully execute the machine check handler in response to the fatal cache error if system logic does not ensure forward progress on the System Bus under this scenario.

Workaround: System logic should ensure completion of the outstanding transactions. Note that during recovery from
a fatal data ECC error, memory image coherency of the BWL with respect to BRL/BRIL transactions is not important. Forward progress is the primary requirement.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P46. L2 Cache May Contain Stale Data in the Exclusive State

Problem: If a cacheline (A) is in Modified (M) state in the write-combining (WC) buffers and in the Invalid (I) state in the L1 cache and it's adjacent sector (B) is in the Invalid (I) state and the following scenario occurs:

1. A read to B misses in the L2 cache and allocates cacheline B and its associated second-sector pre-fetch into an almost full bus queue,
2. A Bus Read Line (BRL) to cacheline B completes with HIT# and fills data in Shared (S) state,
3. The bus queue full condition causes the prefetch to cacheline A to be cancelled, cacheline A will remain M in the WC buffers and I in the L2 while cacheline B will be in the S state. Then, if the further conditions occur:
 1. Cacheline A is evicted from the WC Buffers to the bus queue which is still almost full,
 2. A hardware prefetch Read for Ownership (RFO) to cacheline B, hits the S state in the L2 and observes cacheline A in the I state, allocates both cachelines,



3. An RFO to cacheline A completes before the WC Buffers write modified data back, filling the L2 with stale data,
4. The write-back from the WC Buffers completes leaving stale data, for cacheline A, in the Exclusive (E) state in the L2 cache.

Implication: Stale data may be consumed leading to unpredictable program execution. Intel has not been able to reproduce this erratum in commercial software.

Workaround: It is possible for BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P47. Re-mapping the APIC Base Address to a Value Less Than or Equal to 0xDC001000 may Cause IO and Special Cycle Failure

Problem: Re-mapping the APIC base address from its default can cause conflicts with either I/O or special cycle bus transactions.

Implication: Either I/O or special cycle bus transactions can be redirected to the APIC, instead of appearing on the front-side bus.

Workaround: Use any APIC base addresses above 0xDC001000 as the relocation address.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P48. Erroneous BIST Result Found in EAX Register After Reset

Problem: The processor may show an erroneous BIST (built-in self test) result in the EAX register bit 0 after reset.

Implication: When this erratum occurs, an erroneous BIST failure will be reported in the EAX register bit 0, however this failure can be ignored since it is not accurate.

Workaround: It is possible for BIOS to workaround this issue by masking off bit 0 in the EAX register where BIST results are written.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P49. Processor Does Not Flag #GP on Non-Zero Write to Certain MSRs

Problem: When a non-zero write occurs to the upper 32 bits of IA32_CR_SYSENTER_EIP or IA32_CR_SYSENTER_ESP, the processor should indicate a general protection fault by flagging #GP. Due to this erratum, the processor does not flag #GP.

Implication: The processor unexpectedly does not flag #GP on a non-zero write to the upper 32 bits of

Implication: IA32_CR_SYSENTER_EIP or IA32_CR_SYSENTER_ESP. No known commercially available operating system has been identified to be affected by this erratum.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P50. Simultaneous Assertion of A20M# and INIT# May Result in Incorrect Data Fetch

Problem: If A20M# and INIT# are simultaneously asserted by software, followed by a data access to the 0xFFFFFXXX memory region, with A20M# still asserted, incorrect data will be accessed. With A20M# asserted, an access to 0xFFFFFXXX should result in a load from physical address 0xFFEFFXXX. However, in the case of A20M# and INIT# being asserted together, the data load will actually be from the physical address 0xFFFFFXXX. Code accesses are not affected by this erratum.

Implication: Processor may fetch incorrect data, resulting in BIOS failure.

Workaround: Deasserting and reasserting A20M# prior to the data access will workaround this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*.

P51. Processor Does not Respond to Break Requests From ITP

Problem: On power-up and low-power state transitions, the processor's TAP circuitry may remain in the Tap-Logic-Reset (TLR) state

Workaround: None identified

Implication: The ITP is unable to cause a break on reset in the processor, which may prevent the loading of processor and chipset registers, or affect the ability to debug from cold boot and low power transitions

Status: For the stepping affected, see the *Summary of Changes* at the beginning of this section.

P52. Glitches on Address and Data Strobe Signals May Cause System Shutdown

Problem: When the Processor Signature instruction is executed with EAX = 2 on a processor without Hyper-Threading Technology or with Hyper-Threading Technology disabled via power on configuration, it should return a value of 51h in EAX[15:8] to indicate that the Instruction Translation Lookaside Buffer (ITLB) has 128 entries. On a processor with Hyper-Threading Technology enabled, the processor should return 50h (64 entries). Due to this erratum, the Processor Signature instruction always returns 50h (64 entries).

Implication: Software may incorrectly report the number of ITLB entries. Operation of the processor is not affected.

Workaround: None identified.

Status: For the steppings affected see the Summary of Changes at the beginning of this section.

P53. A write to APIC Task Priority Register (TPR) that lowers priority may seem to have not occurred

Problem: In respect to the retirement of instructions, stores to the uncacheable memory-based APIC register space are handled in a non-synchronized way. For example if an instruction that masks the interrupt flag, e.g. CLI, is executed soon after an uncacheable write to the Task Priority Register (TPR) that lowers the APIC priority, the interrupt masking operation may take effect before the actual priority has been lowered. This may cause interrupts whose priority is lower than the initial TPR, but higher than the final TPR, to not be serviced until the interrupt flag is finally cleared, i.e. by STI instruction. Interrupts will remain pending and are not lost

Implication: This condition may allow interrupts to be accepted by the processor but may delay their service

Workaround This non-synchronization can be avoided by issuing an APIC register read after the APIC register write. This will force the store to the APIC register before any subsequent instructions are executed. No commercial operating system is known to be impacted by this erratum

Status: For the steppings affected see the Summary of Changes at the beginning of this section.

P54. Stop - Clock Assertion May Cause a System to Hang

Problem: Inappropriate assertion of the STPCLK# signal may cause a system to hang. The system may hang upon the following inappropriate assertions of the STPCLK# signal:

1. STPCLK# is asserted prior to the first time a logical processor is a wakened from the wait-for-SIPI State.
2. STPCLK_ACK bus cycle is deferred by an agent external to the processor for a period of time long enough for the chipset to deassert and then reassert STPCLK#. A processor supporting Hyper-Threading Technology may fail to detect the deassertion/reassertion and hence will not generate a STPCLK_ACK in response to the second STPCLK# assertion.

Implication: When this erratum occurs, the system may hang.

Workaround: BIOS should initialize the second thread of the processor supporting Hyper-Threading Technology prior to STPCLK# assertion. Additionally, it is possible for the BIOS to contain a workaround for this erratum

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P55. Store to Load Data Forwarding may Result in Switched Data Bytes

Problem: If in a short window, after an instruction that updates a segment register has executed but not yet retired, there is a load occurring to an address that matches a recent previous store operation but the data size is smaller than the size of the store, the resulting data forwarded from the store to the load may have some of the lower bytes switched

Implication: If this erratum occurs, the processor may execute with incorrect data

Workaround: It is possible for the BIOS to contain a workaround for this erratum

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P56. ITP Cannot Continue Single Step Execution after the First Breakpoint

Problem: ITP will not continue in single step execution after the first software breakpoint. ITP is unable to reset the Resume Flag (RF) bit in the EFLAGS Register

Implication: The processor will break at the instruction breakpoint address instead of single stepping

Workaround: Execution after the break will continue if DR7 bit 1 (Global Breakpoint Enable) is manually cleared.

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P57. Parity Error in the L1 Cache May Cause the Processor to Hang

Problem: If a locked operation accesses a line in the L1 cache that has a parity error, it is possible that the processor may hang while trying to evict the line

Implication: If this erratum occurs, it may result in a system hang. Intel has not observed this erratum with any commercially available software

Workaround: None

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P58. The TCK Input in the Test Access Port (TAP) is Sensitive to Low Clock Edge Rates and Prone to Noise Coupling Onto TCK's Rising or Falling Edges

Problem: TCK is susceptible to double clocking when low amplitude noise occurs on TCK edge, while it is crossing the receiver's transition region. TAP failures tend to increase with increases in background system noise

Implication: This only impacts JTAG/TAP accesses to the processor. Other bus accesses are not affected

Workaround: To minimize the effects of this issue, reduce noise on the TCK-net at the processor relative to ground, and position TCK relative to BCLK to minimize the TAP error rate. Decreasing rise times to under 800ps reduced the failure rate but does not stop all failures

Status: For the steppings affected see the Summary of Changes at the beginning of this section



P59. Disabling a Local APIC Disables Both Logical Processor on a Hyper-Threading Technology Enabled Processor

Problem: Disabling a local APIC on one logical processor of a Hyper-Threading Technology enabled processor by clearing bit 11 of the IA32_APIC_BASE MSR will effectively disable the local APIC on the other logical processor

Implication: Disabling a local APIC on one logical processor prevents the other logical processor from sending or receiving interrupts. Multiprocessor Specification compliant BIOSs and multiprocessor operating systems typically leave all local APICs enabled preventing any end-user visible impact from this erratum

Workaround: Do not disable the local APICs in a Hyper-Threading Technology enabled processor

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P60. Using STPCLK and Executing Code From Very Slow Memory Could Lead to a System Hang

Problem: The system may hang when the following conditions are met:

- 1) Periodic STPCLK mechanism is enabled via the chipset
- 2) Hyper-Threading Technology is enabled
- 3) One logical processor is waiting for an event (i.e. hardware interrupt)
- 4) The other logical processor executes code from very slow memory such that every code fetch is deferred long enough for the STPCLK to be re-asserted

Implication: If this erratum occurs, the processor will go into and out of the sleep state without making forward progress, since the logical processor will not be able to service any pending event. This erratum has not been observed in any commercial platform running commercial software

Workaround: none

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P61. Simultaneous Cacheline Eviction From L2 and L3 Caches may Result in the Write Back of Stale Data

Problem: If a cacheline is evicted simultaneously from both the L2 and L3 caches, and the internal bus queues are full, an older L3 eviction may be allowed to remain in an internal queue entry. If in a narrow timing window an external snoop is generated the data from the older eviction may be used to respond to the external snoop

Implication: In the event that this erratum occurs the contents of memory will be incorrect. This may result in application, operating system or system failure

Workaround: BIOS may contain a workaround for this erratum

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P62. The State of the Resume Flag (RF Flag) in a Task-State Segment (TSS) May be

Problem: After executing a JMP instruction to the next (or other) task through a hardware task switch, it is possible for the state of the RF flag (in the EFLAGS register image) to be incorrect

Implication: The RF flag is normally used for code breakpoint management during debug of an application. It is not typically used during normal program execution. Code breakpoints or single step debug behavior in the presence of hardware task switches, therefore, may be unpredictable as a result of this erratum. This erratum has not been observed in commercially available software

Workaround: none

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P63. Changes to CR3 Register do not Fence Pending Instruction Page

Problem: When software writes to the CR3 register, it is expected that all previous/outstanding code, data accesses and page walks are completed using the previous value in CR3 register. Due to this erratum, it is possible that a pending instruction page walk is still in progress, resulting in an access (to the PDE portion of the page table) that may be directed to an incorrect memory address

Implication: The results of the access to the PDE will not be consumed by the processor so the return of incorrect data is benign. However, the system may hang if the access to the PDE does not complete with data (e.g. infinite number of retries).

Workaround: It is possible for the BIOS to have a workaround for this erratum.

Status: For the steppings affected see the Summary of Changes at the beginning of this section

P64. Simultaneous Page Faults at Similar Page Offsets on Both Logical Processors of an Hyper-Threading Technology Enabled Processor May Cause Application Failure

Problem: An incorrect value of CR2 may be presented to one of the logical processors of an HT Technology enabled processor if a page access fault is encountered on one logical processor in the same clock cycle that the other logical processor also encounters a page fault. Both accesses must cross the same 4 byte aligned offset for this erratum to occur. Only a small percentage of such simultaneous accesses are vulnerable. The vulnerability of the alignment for any given fault is dependent on the state of other circuitry in the processor. Additionally, a third fault from an access that occurs sequentially after one of these simultaneous faults has to be pending at the time of the simultaneous faults.



This erratum is caused by a one-cycle hole in the logic that controls the timing by which a logical processor is allowed to access an internal asynchronous fault address register. The end result is that the value of CR2 presented to one logical processor may be corrupted.

Implication: The operating system is likely to terminate the application that generated an incorrect value of CR2.

Workaround: An operating system or page management software can significantly reduce the already small possibility of encountering this failure by restarting or retrying the faulting instruction and only terminate the application on a subsequent failures of the same instruction. It is possible for BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Tables of Changes*

Specification Changes

The Specification Changes listed in this section apply to the following documents:

- *Intel® Xeon™ Processor at 1.40 GHz, 1.50 GHz, 1.70 and 2 GHz Datasheet* (Order Number 249665)
- *Intel® Xeon™ Processor with 512 KB L2 Cache at 1.80 GHz to 3.0 GHz Datasheet* (Order Number 298642)
- *Intel® Xeon™ Processor with 533 MHz Front Side Bus at 2 GHz to 3.06 GHz Datasheet* (Order Number 252135)
- *Low Voltage Intel® Xeon™ Processor at 1.60GHz to 2.0 GHz Datasheet* (Order Number 273766)
- *IA-32 Intel® Architecture Software Developer's Manual, Volumes 1, 2, and 3* (Order Numbers 245470, 245471, and 245472, respectively)

All Specification Changes will be incorporated into a future version of the appropriate Intel Xeon processor documentation.

P1. Context ID Feature Added to Processor Signature Instruction Feature Flags/IA32_MISC_Enable Registers

IA32_MISC_ENABLE register, bit 24 status has changed from Reserved to the following definition:

IA32_MISC_ENABLE – Miscellaneous Enables Register, bit # 24

MSR Address: 01A0h Accessed as a Qword
 Default Value: High Dword XXXX XXXXh
 Low Dword XXXX XXXX XXXX XXXX XXXX XX00 X0X0 0001b
 Access: Read/Write
 Type: Shared

IA32_MISC_ENABLE is a 64-bit register accessed only when referenced as a Qword through a RDMSR or WRMSR instruction.

Bit 24 of the IA32_MISC_ENABLE status has changed from Reserved to the following:

Bit	Descriptions
24	<p>L1 Data Cache Context Mode (R/W). When set to a '1' this bit places the L1 Data Cache into shared mode. When set to a '0' (default) this bit places the L1 Data Cache into adaptive mode.</p> <p>When this bit is set to a '0', adaptive mode, the Page Directory Base Register contained in CR3 must be identical across all logical processors.</p> <p>Note: If the Context ID feature flag, ECX[10], is not set to a '1' after executing Processor Signature Instruction with EAX = 1, then this feature is not supported and BIOS must not alter the contents of this bit location.</p>



In the Processor Signature Instruction function 1 feature information, bit 10 of ECX register (ECX[10]) has been assigned as a flag to identify “Context ID feature”. The status has changed from Reserved to the following:

ECX [Bits]	Descriptions of Feature Flag Value
10	Context ID. A value of ‘1’ indicates the L1 data cache mode can be set to either adaptive mode or shared mode. A value of ‘0’ this feature is not supported. See definition of the IA32_MISC_ENABLE MSR Bit 24 (L1 Data Cache Context Mode) for more details.

Specification Clarifications

The Specification Clarifications listed in this section apply to the following documents:

- *Intel® Xeon™ Processor at 1.40 GHz, 1.50 GHz, 1.70 and 2 GHz Datasheet (Order Number 249665)*
- *Intel® Xeon™ Processor with 512 KB L2 Cache at 1.80 GHz to 3.0 GHz Datasheet (Order Number 298642)*
- *Intel® Xeon™ Processor with 533 MHz Front Side Bus at 2 GHz to 3.06 GHz Datasheet (Order Number 252135)*
- *Low Voltage Intel® Xeon™ Processor at 1.60GHz to 2.0 GHz Datasheet (Order Number 273766)*
- *IA-32 Intel® Architecture Software Developer's Manual, Volumes 1, 2, and 3 (Order Numbers 245470, 245471, and 245472, respectively)*

All Specification Clarifications will be incorporated into a future version of the appropriate Intel Xeon processor documentation.

P1: Thermal Specification Clarification – Intel Xeon Processor

In the “Thermal Specifications” section of the Intel® Xeon™ Processor at 1.40 GHz, 1.50 GHz, 1.70 GHz and 2 GHz, the thermal specification verbiage has now changed to emphasize the long-term reliability performance degradation if the thermal solutions are under-designed. In order to clarify the purpose of the thermal monitor feature, the description in that section was modified. This change will be incorporated in the next revision of the Intel® Xeon™ Processor at 1.40 GHz, 1.50 GHz, 1.70 GHz and 2 GHz Datasheet. See below for thermal specification clarification changes highlighted in blue.

The Intel® Xeon Processor requires a thermal solution to maintain temperatures within the operating limits as set forth in Section 6.1. Any attempt to operate the processor outside the operating limits may result in permanent damage to the processor and potentially other components within the system. As processor technology pushes to faster speeds and higher performance, thermal management becomes increasingly crucial when building computer systems. Maintaining the proper thermal environment is key to reliable, long-term system operation.

A complete thermal solution includes both component and system level thermal management features. Component level thermal solutions can include active or passive heatsinks attached to the processor Integrated Heat Spreader (IHS). Typical system level thermal solutions may consist of system fans combined with ducting and venting..

This section provides data necessary for developing a complete thermal solution. For more information on designing a component level thermal solution, refer to Intel Xeon Processor Thermal Design Guidelines

6.1 Thermal Specifications



To allow for the optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed such that the processor remains between the minimum and maximum case temperature (TC) specifications when operating at or below the Thermal Design Power (TDP) value listed per frequency in Table 31. Thermal solutions not designed to provide this level of thermal capability may affect the long-term reliability of the processor and system. For more details on thermal solution design, please refer to the appropriate processor thermal design guidelines.

The case temperature is defined at the geometric top center of the processor IHS. Analysis indicates that real applications are unlikely to cause the processor to consume maximum power dissipation for sustained periods of time. Intel recommends that complete thermal solution designs target the Thermal Design Power (TDP) indicated in Table 31 instead of the maximum processor power consumption. The Thermal Monitor feature is intended to help protect the processor in the unlikely event that an application exceeds the TDP recommendation for a sustained period of time. For more details on the usage of this feature, refer to Section 7.3. To ensure maximum flexibility for future requirements, systems should be designed to the Flexible Motherboard Guidelines (FMB), even if a processor with a lower thermal dissipation is currently planned. In all cases, the Thermal Monitor feature must be enabled for the processor to remain within specification.

Section 6.2.1 Processor Case Temperature Measurement

The maximum and minimum case temperature (TCASE) for the Intel Xeon Processor is specified in Table 31. This temperature specification is meant to help ensure proper operation of the processor. Figure 31 illustrates where Intel recommends TCASE thermal measurements should be made.

Section 7.0 Features

Section 7.3 Thermal Monitor

The Thermal Monitor feature helps control the processor temperature by activating the Thermal Control Circuit (TCC) when the processor silicon reaches its maximum operating temperature. The TCC reduces processor power consumption by modulating (starting and stopping) the internal processor core clocks. The Thermal Monitor feature must be enabled for the processor to be operating within specifications. The temperature at which Thermal Monitor activates the thermal control circuit is not user configurable and is not software visible. Bus traffic is snooped in the normal manner, and interrupt requests are latched (and serviced during the time that the clocks are on) while the TCC is active.

When the Thermal Monitor feature is enabled, and a high temperature situation exists (i.e. TCC is active), the clocks will be modulated by alternately turning the clocks off and on at a duty cycle specific to the processor (typically 30-50%). Clocks often will not be off for more than 3.0 microseconds when the TCC is active. Cycle times are processor speed dependent and will decrease as processor core frequencies increase. A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the TCC goes inactive and clock modulation ceases.

With a properly designed and characterized thermal solution, it is anticipated that the TCC would only be activated for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss, and in some cases may result in a **TC** that exceeds the specified maximum temperature and may affect the long-term reliability of the processor. In addition, a thermal solution that is significantly under-designed may not be capable of cooling the processor even when the TCC is active continuously. Refer to the Intel Xeon Processor Thermal Design Guidelines for information on designing a thermal solution.

The duty cycle for the TCC, when activated by the Thermal Monitor, is factory configured and cannot be modified. The Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines.

P2: Thermal Specification Clarification – Intel Xeon Processor with 512-KB L2 cache

In the “Thermal Specifications” section of the Intel® Xeon™ Processor with 512-KB L2 Cache at 1.80 GHz to 3 GHz Datasheet, the thermal specification verbiage has now changed to emphasize the long-term reliability performance degradation if the thermal solutions are under-designed. In order to clarify the purpose of the thermal monitor feature, the description in that section was modified. This change will be incorporated in the next revision of the Intel® Xeon™ Processor with 512-KB L2 Cache at 1.80 GHz to 3 GHz Datasheet See below for thermal specification clarification changes highlighted in blue.

The Intel® Xeon Processor requires a thermal solution to maintain temperatures within the operating limits as set forth in Section 6.1. Any attempt to operate the processor outside the operating limits may result in permanent damage to the processor and potentially other components within the system. As processor technology pushes to faster speeds and higher performance, thermal management becomes increasingly crucial when building computer systems. Maintaining the proper thermal environment is key to reliable, long-term system operation.

A complete thermal solution includes both component and system level thermal management features. Component level thermal solutions can include active or passive heatsinks attached to the processor Integrated Heat Spreader (IHS). Typical system level thermal solutions may consist of system fans combined with ducting and venting..

This section provides data necessary for developing a complete thermal solution. For more information on designing a component level thermal solution, refer to Intel Xeon Processor Thermal Design Guidelines

6.1 Thermal Specifications

To allow for the optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed such that the processor remains between the minimum and maximum case temperature (TC) specifications when operating at or below the Thermal Design Power (TDP) value listed per frequency in Table 42. Thermal solutions not designed to provide this level of thermal capability may affect the long-term reliability of the processor and system. For more details on thermal solution design, please refer to the appropriate processor thermal design guidelines.

The case temperature is defined at the geometric top center of the processor IHS. Analysis indicates that real applications are unlikely to cause the processor to consume maximum power dissipation for sustained periods of time. Intel recommends that complete thermal solution designs target the Thermal Design Power (TDP) indicated in Table 42 instead of the maximum processor power consumption. The Thermal Monitor feature is intended to help protect the processor in the unlikely event that an application exceeds the TDP recommendation for a sustained period of time. For more details on the usage of this feature, refer to Section 7.3. To ensure maximum flexibility for future requirements, systems should be designed to the Flexible



Motherboard Guidelines (FMB), even if a processor with a lower thermal dissipation is currently planned. In all cases, the Thermal Monitor feature must be enabled for the processor to remain within specification.

Section 6.2.1 Processor Case Temperature Measurement

The maximum and minimum case temperature (TCASE) for the Intel Xeon Processor is specified in Table 42. This temperature specification is meant to help ensure proper operation of the processor. Figure 40 illustrates where Intel recommends TCASE thermal measurements should be made.

Section 7.0 Features

Section 7.3 Thermal Monitor

The Thermal Monitor feature helps control the processor temperature by activating the Thermal Control Circuit (TCC) when the processor silicon reaches its maximum operating temperature. The TCC reduces processor power consumption by modulating (starting and stopping) the internal processor core clocks. The Thermal Monitor feature must be enabled for the processor to be operating within specifications. The temperature at which Thermal Monitor activates the thermal control circuit is not user configurable and is not software visible. Bus traffic is snooped in the normal manner, and interrupt requests are latched (and serviced during the time that the clocks are on) while the TCC is active.

When the Thermal Monitor feature is enabled, and a high temperature situation exists (i.e. TCC is active), the clocks will be modulated by alternately turning the clocks off and on at a duty cycle specific to the processor (typically 30-50%). Clocks often will not be off for more than 3.0 microseconds when the TCC is active. Cycle times are processor speed dependent and will decrease as processor core frequencies increase. A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the TCC goes inactive and clock modulation ceases.

With a properly designed and characterized thermal solution, it is anticipated that the TCC would only be activated for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss, and in some cases may result in a **TC** that exceeds the specified maximum temperature and may affect the long-term reliability of the processor. In addition, a thermal solution that is significantly under-designed may not be capable of cooling the processor even when the TCC is active continuously. Refer to the Intel Xeon Processor Thermal Design Guidelines for information on designing a thermal solution.

The duty cycle for the TCC, when activated by the Thermal Monitor, is factory configured and cannot be modified. The Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines.

P4: Thermal Specification Clarification – Low Voltage Intel Xeon Processor

In the “Thermal Specifications” section of the Low Voltage Intel® Xeon™ Processor at 1.6 GHz and 2.0 GHz Datasheet, the thermal specification verbiage has now changed to emphasize the long-term reliability performance degradation if the thermal solutions are under-designed. In order to clarify the purpose of the thermal monitor feature, the description in that section was modified. This change will be incorporated in the next revision of the Low Voltage Intel® Xeon™ Processor at 1.6 GHz and 2.0 GHz Datasheet, See below for thermal specification clarification changes highlighted in blue.

The Low Voltage Intel® Xeon Processor requires a thermal solution to maintain temperatures within the operating limits as set forth in Section 6.1. Any attempt to operate the processor outside the operating limits may result in permanent damage to the processor and potentially other components within the system. As processor technology pushes to faster speeds and higher performance, thermal management becomes increasingly crucial when building computer systems. Maintaining the proper thermal environment is key to reliable, long-term system operation.

A complete thermal solution includes both component and system level thermal management features. Component level thermal solutions can include active or passive heatsinks attached to the processor Integrated Heat Spreader (IHS). Typical system level thermal solutions may consist of system fans combined with ducting and venting..

This section provides data necessary for developing a complete thermal solution. For more information on designing a component level thermal solution, refer to Intel Xeon Processor Thermal Design Guidelines

6.1 Thermal Specifications

To allow for the optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed such that the processor remains between the minimum and maximum case temperature (TC) specifications when operating at or below the Thermal Design Power (TDP) value listed per frequency in Table 32. Thermal solutions not designed to provide this level of thermal capability may affect the long-term reliability of the processor and system. For more details on thermal solution design, please refer to the appropriate processor thermal design guidelines.

The case temperature is defined at the geometric top center of the processor IHS. Analysis indicates that real applications are unlikely to cause the processor to consume maximum power dissipation for sustained periods of time. Intel recommends that complete thermal solution designs target the Thermal Design Power (TDP) indicated in Table 32 instead of the maximum processor power consumption. The Thermal Monitor feature is intended to help protect the processor in the unlikely event that an application exceeds the TDP recommendation for a sustained period of time. For more details on the usage of this feature, refer to Section 7.3. To ensure maximum flexibility for future requirements, systems should be designed to the Flexible Motherboard Guidelines (FMB), even if a processor with a lower thermal dissipation is currently planned. In all cases, the Thermal Monitor feature must be enabled for the processor to remain within specification.

Section 6.2.1 Processor Case Temperature Measurement



The maximum and minimum case temperature (TCASE) for the Intel Xeon Processor is specified in Table 32. This temperature specification is meant to help ensure proper operation of the processor. Figure 34 illustrates where Intel recommends TCASE thermal measurements should be made.

Section 7.0 Features

Section 7.3 Thermal Monitor

The Thermal Monitor feature helps control the processor temperature by activating the Thermal Control Circuit (TCC) when the processor silicon reaches its maximum operating temperature. The TCC reduces processor power consumption by modulating (starting and stopping) the internal processor core clocks. The Thermal Monitor feature must be enabled for the processor to be operating within specifications. The temperature at which Thermal Monitor activates the thermal control circuit is not user configurable and is not software visible. Bus traffic is snooped in the normal manner, and interrupt requests are latched (and serviced during the time that the clocks are on) while the TCC is active.

When the Thermal Monitor feature is enabled, and a high temperature situation exists (i.e. TCC is active), the clocks will be modulated by alternately turning the clocks off and on at a duty cycle specific to the processor (typically 30-50%). Clocks often will not be off for more than 3.0 microseconds when the TCC is active. Cycle times are processor speed dependent and will decrease as processor core frequencies increase. A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the TCC goes inactive and clock modulation ceases.

With a properly designed and characterized thermal solution, it is anticipated that the TCC would only be activated for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss, and in some cases may result in a **TC** that exceeds the specified maximum temperature and may affect the long-term reliability of the processor. In addition, a thermal solution that is significantly under-designed may not be capable of cooling the processor even when the TCC is active continuously. Refer to the Intel Xeon Processor Thermal Design Guidelines for information on designing a thermal solution.

The duty cycle for the TCC, when activated by the Thermal Monitor, is factory configured and cannot be modified. The Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines.